The Revelations of Acoustic Waves

Livermore researchers are developing advanced techniques to extract information from acoustic signals.

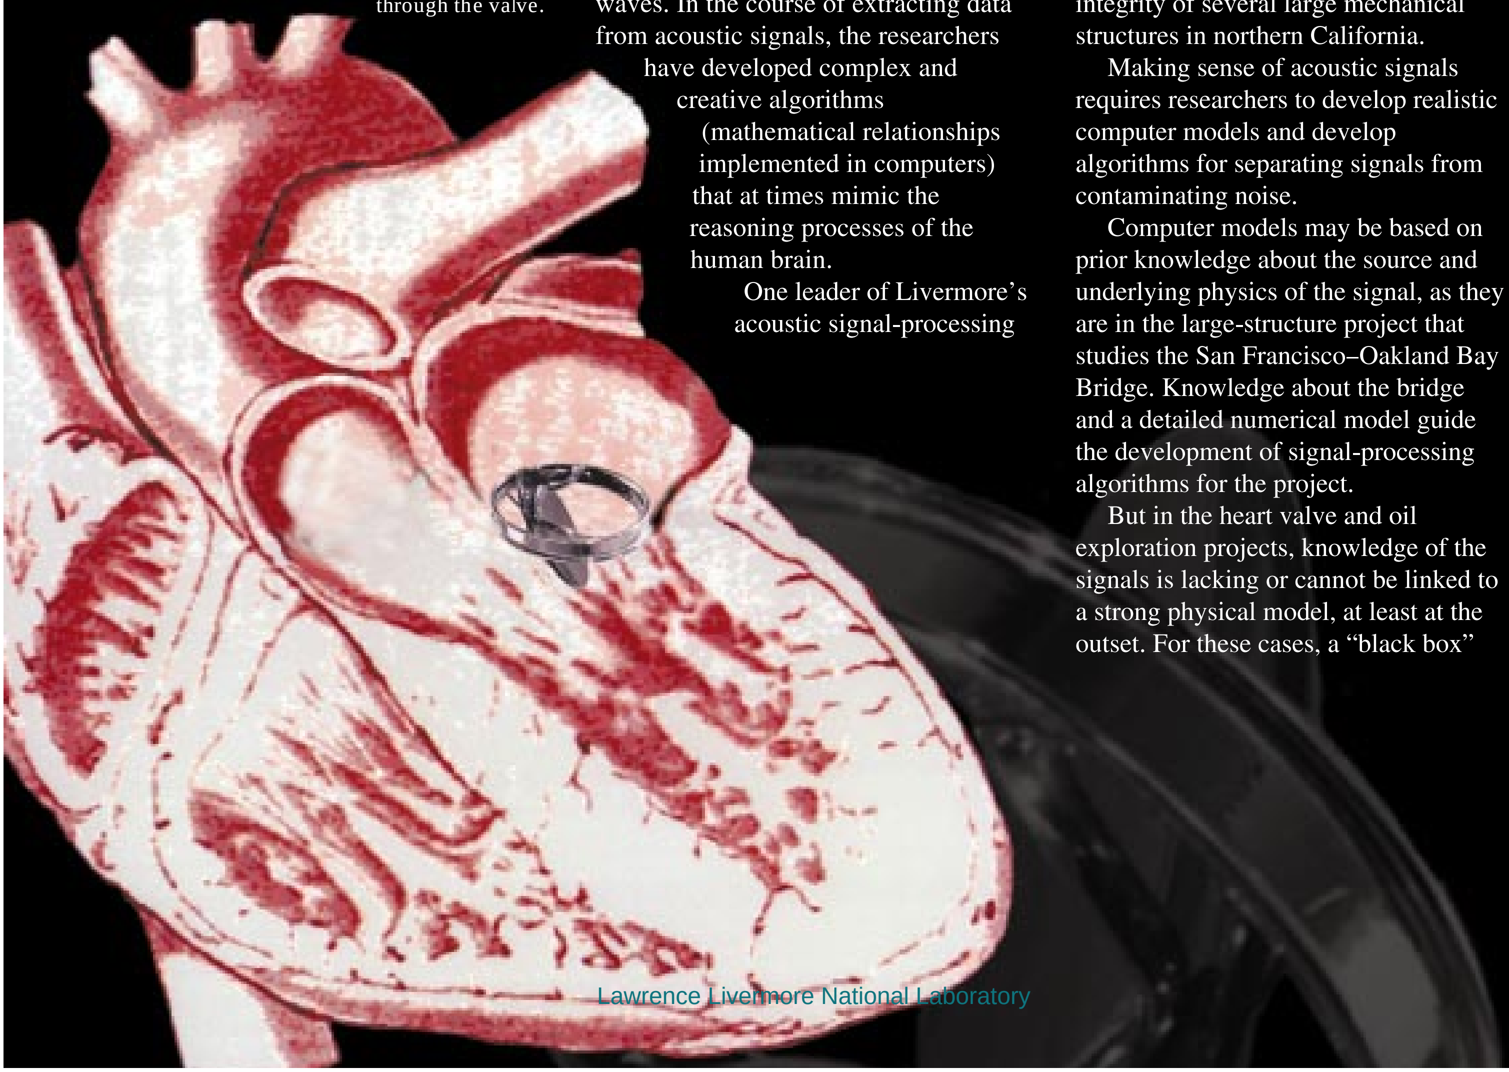

Figure 1. A heart with an artificial heart valve. The valve consists of a disk held in place by two struts that let it flip open and shut during pumping of blood through the valve.

FROM analyzing speech to recording earthquakes, tracking submarines, or imaging a fetus, measuring and analyzing acoustic signals are increasingly important in modern society. Acoustic waves are simply disturbances involving mechanical vibrations in solids, liquids, or gases. Lawrence Livermore researchers are developing advanced techniques for extracting and interpreting the information in these waves. In the course of extracting data from acoustic signals, the researchers have developed complex and creative algorithms (mathematical relationships implemented in computers) that at times mimic the reasoning processes of the human brain.

One leader of Livermore's acoustic signal-processing research is electronics engineer Greg Clark, who is involved in three disparate acoustics projects: heart valve classification, where acoustic signal processing is determining whether an artificial heart valve is intact or needs replacing; oil exploration, where Livermore experts are automating a key procedure used for locating undersea oil deposits; and large-structure analysis, where Livermore is preparing to use acoustic wave vibrations to assess the integrity of several large mechanical structures in northern California.

Making sense of acoustic signals requires researchers to develop realistic computer models and develop algorithms for separating signals from contaminating noise.

Computer models may be based on prior knowledge about the source and underlying physics of the signal, as they are in the large-structure project that studies the San Francisco–Oakland Bay Bridge. Knowledge about the bridge and a detailed numerical model guide the development of signal-processing algorithms for the project.

But in the heart valve and oil exploration projects, knowledge of the signals is lacking or cannot be linked to a strong physical model, at least at the outset. For these cases, a "black box" model, derived only from the data (that is, input and output signals), without details of the underlying physics, is used to guide the development of algorithms.

For the three projects, Clark uses advanced signal-processing techniques, including statistical neural networks, which are systems of computer programs that approximate the operation of the human brain. Current uses for neural networks include predicting weather patterns, interpreting nucleotide sequences, and recognizing features in images. In a supervised learning mode, a neural network is "trained" with large numbers of examples and rules about data relationships. This training endows the network with the ability to make reasonable yes–no decisions on whether, for example, a geologic data plot indicates a geologic layer that could mark the presence of an oil deposit, or whether energy in a frequency spectrum of a recording signifies a damaged artificial heart valve.

In every acoustic signal project, vital data must be separated from noise that contaminates and inevitably degrades signal quality. The noise is caused both by the surrounding environment and the very system recording the signals. For example, the remote system designed by Livermore engineers for monitoring large structures can introduce noise into the signal in the course of relaying it over cellular phone lines to the Laboratory for analysis. In another example, the delicate sounds of a heart valve flipping up and down can be buried by acoustic scattering inside the body. Similarly, multiple ricocheting reflections from underwater explosions can contaminate the precise data needed to isolate geologic strata.

Often, Livermore engineers can reduce noise by using filters for certain frequency spectra. For example, if they know that the structural failure of a bridge will cause a vibrational response at 5 hertz (cycles per second), they can design a filter to monitor the energy content in that part of the frequency spectrum.

Hearing the Heart

For medical diagnostic applications, acoustic signals provide advantages in being noninvasive and harmless. These advantages are being exploited by a team of Livermore engineers who are developing an acoustic processing technique for sifting through a seeming cacophony of heart and body sounds to isolate the few telltale signals of a faulty artificial heart valve. Their technique would spare patients, many of them elderly, from open-heart surgery to determine if an artificial valve needs replacement.

The four valves of the human heart continually open and close, allowing blood to be pumped through the heart's four chambers. When a valve becomes diseased, pumping ability decreases. Prosthetic heart valves correct this deficiency and extend the life spans of many people with serious heart conditions. But Livermore engineer and project co-principal investigator Jim Candy points out that prosthetic valves, while extremely reliable, are eventually susceptible to long-term fatigue and structural failure, as might be expected from any mechanical device operating over a long time.

The Livermore experts are working to find ways to identify faulty heart valves made by one medical device manufacturer whose heart valves were implanted from 1979 to 1986 in 86,000 patients. To date, more than 600 of these valves have failed, and more than 300 people have died. A court-appointed panel is funding research to find the best screening technique or combination of techniques to determine, with a high degree of accuracy, if one of the manufacturer's implanted valves is failing. Livermore's acoustical processing method is a leading technique. Tests show it is more effective than x-ray procedures, which cannot capture a clear image of moving heart valves.

An artificial heart valve is essentially a small ring with two small struts welded to it, opposite each other. One is an inlet strut, and the other, an outlet strut. Each is two-legged. The two struts hold a disk that flips up and down to open and close the valve (Figure 1).

Over time, the struts can develop cracks at or near their weld joints and break loose from the ring. Livermore signal experts are trying to detect heart valve failure indicated by one leg of a strut breaking loose from the ring. (If both legs break loose, the valve loses control of blood flow to the heart, leading to death in two of three cases.)

To determine the condition of an artificial valve, the Livermore researchers use acoustic signals of recipients' heart valve sounds that have been recorded at clinics using high-sensitivity microphones. They first collect a database of heart valve sounds of about 100 beats per patient, discarding heartbeats with too low a signal-to-noise ratio or with characteristics statistically different from the other beats. The team uses algorithms to classify the recordings as indicative of a valve that is either intact or not. The algorithms scrutinize the frequency spectra of each opening sound of a valve and select key features of the spectra, usually parts of certain peaks. A statistical pattern classifier—in this case, an artificial neural network trained on the recorded acoustic fingerprint of known faulty valves—decides whether the valve is indeed damaged (Figure 2).

The team's efforts are focused on the sounds made when the valve opens—sounds caused by the disk hitting the outlet strut—because those sounds yield direct information on the condition of

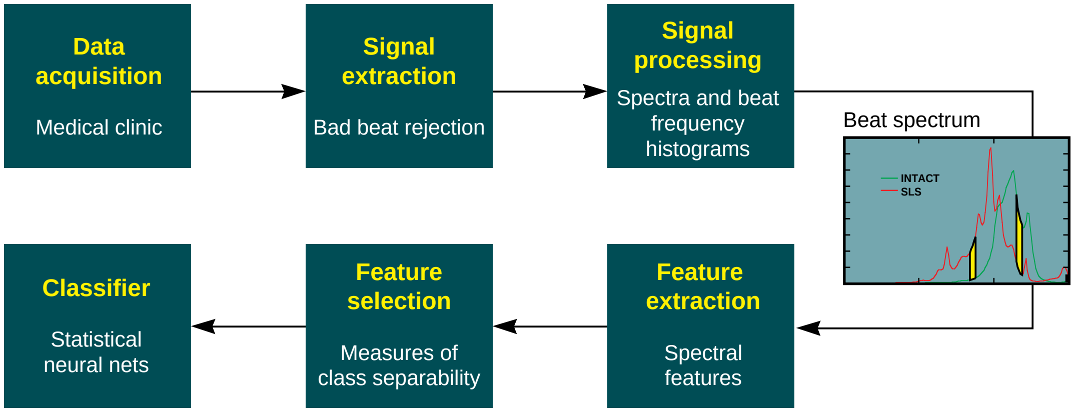

Figure 2. The process for classifying heart valves includes acquiring clinical data, rejecting heartbeats that have low signal-to-noise ratios or are statistical outliers, estimating signal spectra, extracting the spectral features to be used to discriminate between intact and nonintact valves, selecting only the most important features, and finally, classifying the valves as intact or not.

the strut. A closing valve, in contrast, causes the entire ring to vibrate, and that masks the strut's vibrations. Clark compares the sound of a faulty valve to the thud of a cracked bell.

Unfortunately, the opening sounds have much lower intrinsic signal levels than the valve's closing sounds. Candy notes that measuring heart sounds noninvasively in this noisy environment puts significant demands on the signal-processing techniques to extract the desired signals, especially when the data of interest last only 10 to 20 milliseconds. "Finding the opening sound caused by a strut separating from the heart valve is more difficult than searching for a single violin string that is out of tune in an entire orchestra," Candy says.

Clark points out that every valve recording is distorted by the body cavity and the recording process itself. That is why Livermore researchers are conducting studies at a U.S. Navy laboratory in San Diego in which acoustic sensors are submerged in water while collecting the sounds of valves that have been surgically removed from patients. This submersion isolates the pure sounds of both intact and damaged valves. "The test will allow us to measure the heart valve sounds without the acoustic scattering effects caused by the body," Clark says. "The pure data should allow us to mitigate the distortion in patient recordings."

Listening for Oil

Some of the same approaches taken to analyze the subtle sounds of human heart valves are used to locate the interfaces of undersea geologic layers, particularly slate and sandstone layers where oil tends to accumulate around salt domes, or "plugs" (Figure 3). Mapping these geologic layer interfaces (called event horizons or, simply, events) helps geologists decide where to site oil drilling rigs.

One such mapping project is sponsored by the Department of Energy National Gas and Oil Technology Partnership. It is a cooperative effort involving Shell Oil of Houston, Texas, as Lawrence Livermore's industrial partner and a Ph.D. student from the University of California at Davis who is being supported by the Laboratory to work on the project.

The purpose of the Livermore–Shell–UC Davis project is to automate a technique used to analyze acoustic oil

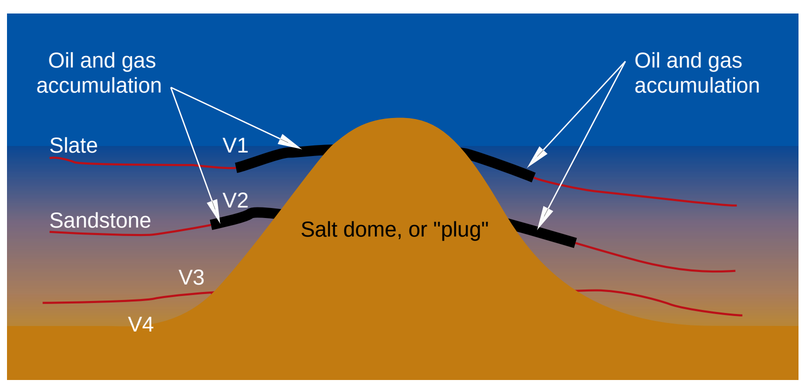

Figure 3. Acoustic techniques can be used to locate the interfaces between undersea geologic layers (usually slate and sandstone) and salt domes, where oil tends to accumulate.

exploration data. The current technique is a significant bottleneck because it is performed manually, at great cost in time and money. The project's goal is to reduce manual effort to only about 0.1 percent of the data processed. Current results show that this goal is achievable, says Clark.

Oil companies obtain acoustic signals by having a ship set off underwater explosions that generate acoustic waves that reflect from geologic layers as deep as 15 to 20 miles under the sea. A 5-mile-long linear array of some 100 hydrophones, which is towed by the ship, measures the signals (Figure 4).

The signals are organized into two-dimensional images, called common reflection point (CRP) panels, which represent vertical sections of the earth (Figure 5a). The panels show multiple reflections of the acoustic waves as they bounce off various strata of earth (Figure 5b). "You get a horrendous number of reflections," says Clark. "It gets very messy to sort out."

Sorting out the reflections depends on correcting for the velocity of the sound waves as they travel through different layers of the earth. If the velocity computer model is correct, the imaged events appear as approximately straight horizontal (flat) lines in the CRP panel, because the true depth of the event horizon is approximately constant across the panel.

If, as typically occurs, the initial velocity computer model is incorrect, the event depths vary across the panel and do not appear flat. As part of an iterative velocity estimation process, an expert must visually inspect the panels to pick out event locations manually. The expert's picks are then used as input to refine the velocity model. This process is repeated several times, until the model produces events imaged as flat lines. The corrected panels are combined to obtain a two-dimensional image of the subsurface strata to help geologists determine where to site an offshore drilling platform.

When Clark inspected the CRP panels at Shell facilities in Houston, he suggested a new approach for analyzing reflection data. Instead of analyzing one signal or a few signal traces at a time, as is conventionally done, he proposed treating the set of 45 traces that forms a CRP panel as a single image. "I pointed out that if you treat the panel as an image, there's a whole set of literature and a lot of powerful tools available to you," he said (Figure 5c).

The Livermore team developed a technique that breaks the panels into small pixels (picture elements) to determine if the data represented within each pixel are part of an event or simply background noise. The technique uses advanced algorithms from the areas of automatic target recognition, computer vision, and signal–image processing. For example, using algorithms similar to those employed for computerized military target recognition, the technique

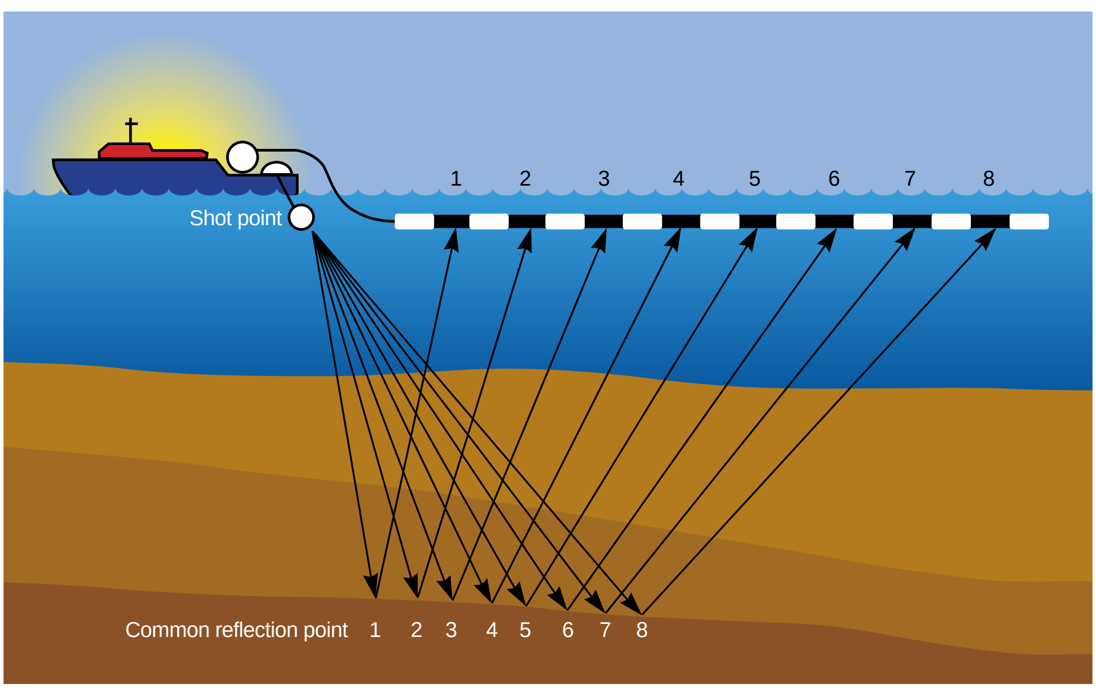

Figure 4. Underwater explosions are set off to generate acoustic waves that reflect off geologic layers and are received by a linear array of equally spaced hydrophones.

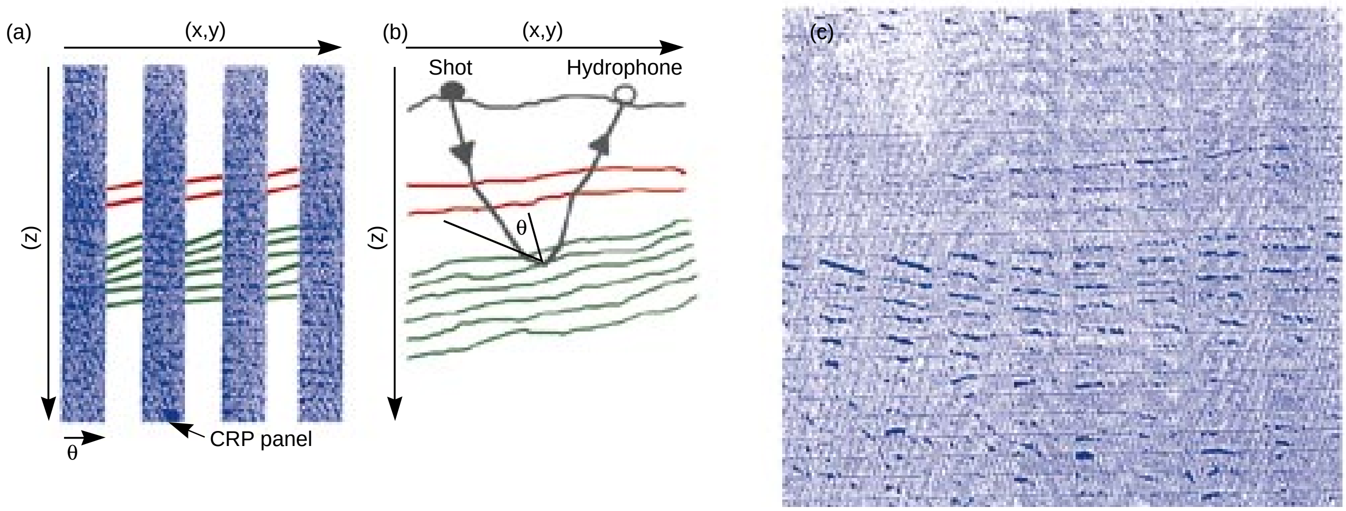

Figure 5. (a) A series of common reflection point (CRP) panels that represent vertical sections of the earth. (b) A depiction of explosive shots causing the wave reflections that are recorded by hydrophones and imaged into CRPs. (c) The CRP panels plotted side by side to form a mosaic.

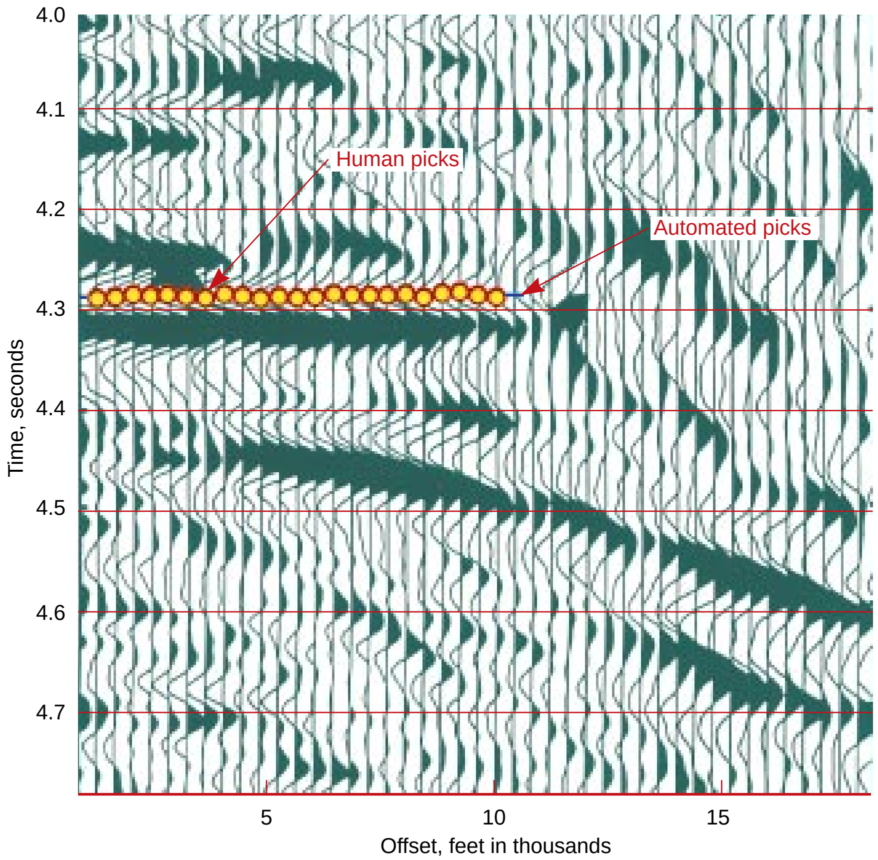

Figure 6. Results of Livermore's automated process for picking out geologic layer interfaces (called events) compare favorably with those attained by experts. The circles depict manual picks, while the blue line, barely visible behind the circles, depicts the automated picks. The manual picks and automated picks overlap so closely they are hard to differentiate.

scrutinizes the neighborhood of each pixel at various orientations and scales and looks for common features with neighboring pixels. Whenever possible, prior knowledge of known events is incorporated into the procedure. The results of the Livermore process compare favorably with those attained by experts (Figure 6).

Clark has made many trips to Shell to collaborate with their scientists on the project. "It's been a wonderful marriage of disciplines," he says. "I introduced them to a lot of advanced signal-processing techniques, and they educated me about their world of exploration geophysics." Clark's Ph.D. research involved analysis of seismograms for Livermore's Comprehensive Test Ban Program, so he already had an elementary geophysical background.

Shell Oil estimates that the Livermore work will significantly affect the oil industry. Shell's current implementation of the Livermore algorithms already reduces data-picking time from about one work day to about 90 minutes. Once the software is completely converted from research code to production code, Shell estimates the cost of performing a single velocity analysis could be reduced from $75,000 and 12 weeks to $6,000 and one week. An oil company performing 100 velocity analyses per year could save nearly $7 million annually. Potential annual savings for the U.S. oil industry could amount to roughly $140 million.

Clark points out that it costs about $1 billion to erect an oil platform in the Gulf of Mexico. Such costs make it crucial to respond quickly to business opportunities. The Livermore signal-processing techniques, providing time savings of a factor of 12, could significantly enhance the industry's responsiveness.

Vibrational Fingerprints

Acoustic signal processing may also make it possible to analyze vibrations and thus assess large mechanical structures for damage after earthquakes or other destructive events. One Livermore project is combining signal processing with advanced numerical models and new remote monitoring systems to better understand large structures and provide a unique way to quickly monitor them for damage.

"Our task is similar to that of the heart valve project: use sophisticated signal processing to enhance our understanding of the way large structures vibrate to find out if there is damage," says project leader and mechanical engineer David McCallen.

There is a critical need for a speedy method to assess the integrity of a structure after an extreme event, adds McCallen, who has worked with Caltrans (California's Department of Transportation) on earthquake-related projects. Current procedures require lengthy, largely visual inspections.

McCallen readily acknowledges the technical challenges of using vibration measurements to determine the health of a structure. "We're asking a lot of our signal-processing people," he says. "We're telling them, we'll give you enough data and insight into the structure so it won't be a black box situation, and we want you to tell us if there is damage, what it is, and where it is."

The Livermore project involves three northern California case studies: the Bixby Creek Bridge in Big Sur, the San Francisco–Oakland Bay Bridge, and the National Ignition Facility (NIF), the world's largest laser under construction at Lawrence Livermore. Each structure previously has been studied at Livermore; as a result, a detailed numerical model exists for each (Figure 7).

The numerical models provide information useful for designing sensors (accelerometers) that Livermore researchers will install on large structures for remote monitoring. The models indicate what frequencies are of interest and also

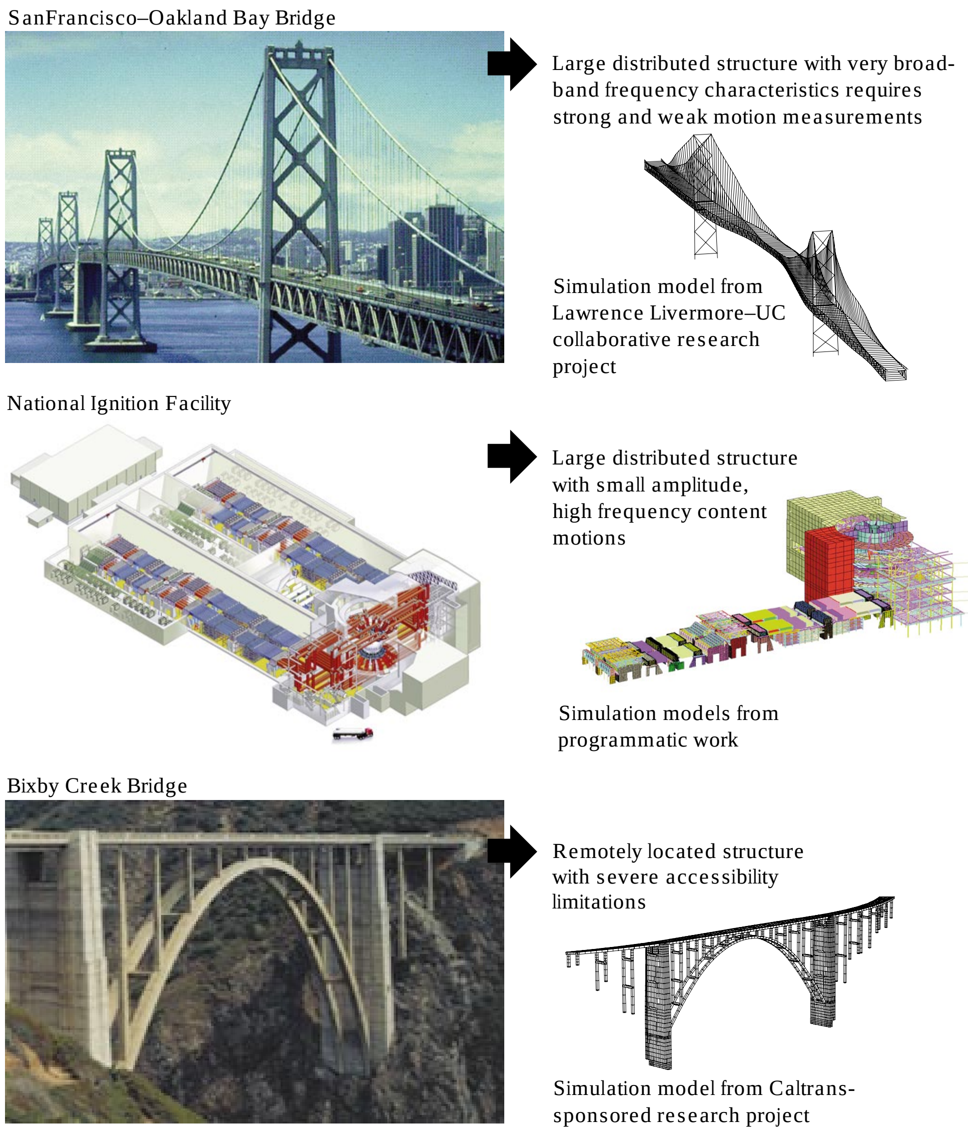

Figure 7. Livermore's detailed numerical models of three large structures are the basis for developing algorithms that analyze acoustic wave vibrations to determine whether these structures have sustained damage in an earthquake or other extreme event.

help determine the best locations for the sensors.

The monitoring system developed by Livermore engineers will continuously record, time stamp, and store sensor data. Researchers can contact the system at any time by cellular phone to download the data for analysis.

In developing signal-processing algorithms, Clark and colleagues integrate data from the sensors, models of sensor noise and the structure's unique environment, and a state–space numerical model. State–space models (common in electronics engineering) transform standard finite-element computer models (common in mechanical engineering) to render them more suitable for signal processing.

The resulting algorithms compare numerical model simulations with measurements of the real structure. A discrepancy between the two is a sign that the structure has suffered damage. Future algorithms will attempt to determine the source of discrepancies between the numerical model and the structure. Statistical classifiers, including artificial neural networks, could then come into play to classify what kind of damage the structure has sustained.

Data were collected recently to test the ability of signal-processing algorithms to detect differences between the numerical model and the actual structure. An experimental structure at the Nevada Test Site, a scale model of a five-story building some 14 feet tall, was the testbed for verifying these algorithms (Figure 8).

Engineers used the experimental structure both to evaluate the sensor systems and to acquire data for

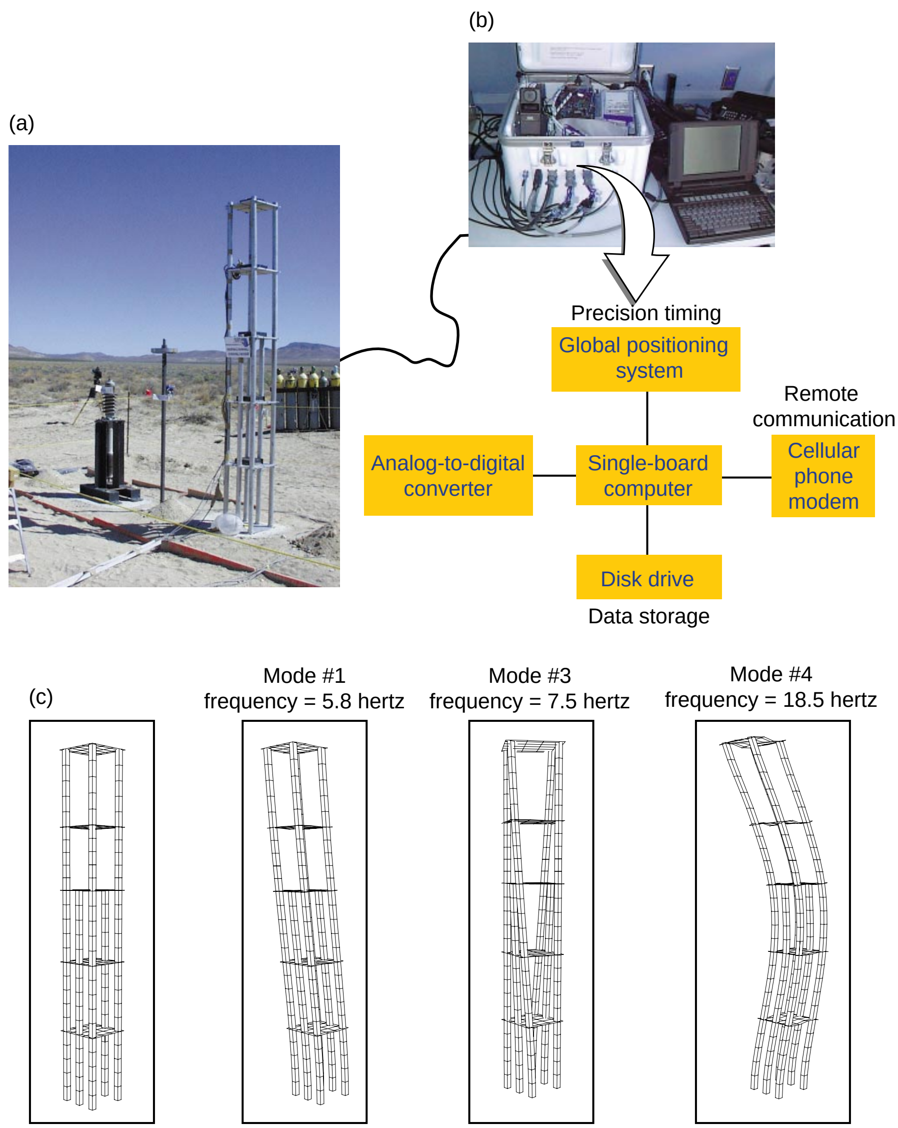

Figure 8. (a) A scale-model building at the Nevada Test Site was used to evaluate how well signal-processing algorithms could detect damage from earthquakes or other events. (b) Vibration data could be analyzed at the experimental site or downloaded in near real-time to remote locations via a cellular phone included in the data acquisition system. (c) Measurements of the scale-model building were compared with finite-element models to determine whether the structure had suffered damage as a result of the simulated earthquake. The figure indicates the first three natural modes computed with the computational model.

evaluating the signal-processing algorithms. For the latter purpose, they simulated an earthquake using a small amount of explosives contained in a rubber bladder. They also vibrated the structure continuously to excite every frequency at which it might vibrate.

This summer, a dozen sensors will be placed on the Bixby Creek Bridge to record vibrational ground motion from small earthquakes. Placement and design of the sensors were guided by the existing numerical model of the bridge, made by Livermore engineers for Caltrans to evaluate the retrofitted bridge in a large earthquake.

Also this summer, a remotely monitored sensor and data acquisition system will be installed at various locations on the San Francisco–Oakland Bay Bridge. The system will monitor both ambient vibrations from traffic and wind and the structure's response to ground motion from small earthquakes. The guiding numerical model of the bridge is a product of a Livermore–University of California at Berkeley project (see *S&TR*, December 1998, p. 18). Data from the system will allow researchers to identify the bridge's "healthy fingerprint" and assess how well the signal-processing tools detect and identify discrepancies between model simulations and measured structures.

Prototype instruments are also being placed on NIF for long-term structural monitoring. The monitoring will help ensure that the giant laser facility's sensitive optical systems can perform under ambient vibration conditions, which include traffic, air conditioners, and other "cultural noise" effects, as well as microearthquakes. The NIF numerical model was made prior to beginning construction of the $1.2-billion laser facility.

The Livermore team hopes the work will provide a structural monitoring capability for Caltrans that can also be applied to critical DOE sites such as hazardous material facilities. In this way, says McCallen, authorities could have a much better handle on assessing damage to important structures and determining response and upgrade priorities.

Listening to the World

Acoustic waves permeate the natural and cultural world. The sound of a heart valve, the acoustic reflection from a pool of oil, and the vibration of a building are only three examples. But techniques developed for these research projects may well be applicable to other research fields. "Everything we're currently developing will help us to solve problems in other areas," says Clark.

—Arnie Heller

Key Words: accelerometers, acoustic signals, algorithms, artificial neural networks, Bixby Creek Bridge, Caltrans, Department of Energy National Gas and Oil Technology Partnership, heart valves, National Ignition Facility, Nevada Test Site, numerical models, oil exploration, San Francisco–Oakland Bay Bridge, signal processing, structure analysis.

For further information contact Greg Clark (925) 423-9759 (clark9@llnl.gov).

About the Scientist

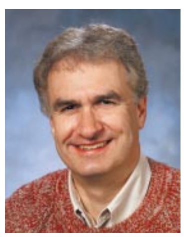

GREGORY A. CLARK received his B.S. and M.S. in electrical engineering from Purdue University in 1972 and 1974, respectively, and his Ph.D. in electrical and computer engineering from the University of California at Santa Barbara in 1981. His research activities are in the theory and application of automatic target recognition, computer vision, sensor fusion, pattern recognition–neural computing, estimation–detection, signal and image processing, and automatic control. He joined Lawrence Livermore in 1974 and is currently a principal investigator for several projects in the Defense Sciences Engineering Division. He has contributed to over a hundred technical publications and serves as a reviewer for several professional journals.